

EV355228095

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

System and Method for Hashing Digital Images

Inventor(s):
Ramarathnam Venkatesan
William Koon

ATTORNEY'S DOCKET NO. MS1-422USC1

RELATED APPLICATIONS

This is a continuation of pending U.S. patent application serial number 09/421,986, entitled "System and Method for Hashing Digital Images," filed on October 19, 1999.

TECHNICAL FIELD

This invention relates to systems and methods for hashing digital bit streams such as digital images. This invention further relates to database systems and methods that utilize the hashing techniques for indexing bit streams and protecting copyrights in the bit streams.

BACKGROUND

Digital images offer many advantages over conventional media in terms of image quality and ease of transmission. However, digital images consume large amounts of memory space. With the ever increasing popularity of the Internet, digital images have become a mainstay ingredient of the Web experience, buoyed by such advances as the increasing speed at which data is carried over the Internet and improvements in browser technology for rendering such images. Everyday, numerous digital images are added to Web sites around the world.

As image databases grow, the needs for indexing them and protecting copyrights in the images are becoming increasingly important. The next generation of database management software will need to accommodate solutions for fast and efficient indexing of digital images and protection of copyrights in those digital images.

1 A hash function is one probable solution to the image indexing and
2 copyright protection problem. Hash functions are used in many areas such as
3 database management, querying, cryptography, and many other fields involving
4 large amounts of raw data. A hash function maps large unstructured raw data into
5 relatively short, structured identifiers (the identifiers are also referred to as “hash
6 values” or simply “hash”). By introducing structure and order into raw data, the
7 hash function drastically reduces the size of the raw data into short identifiers. It
8 simplifies many data management issues and reduces the computational resources
9 needed for accessing large databases.

10 Thus, one property of a good hash function is the ability to produce small-
11 size hash values. Searching and sorting can be done much more efficiently on
12 smaller identifiers as compared to the large raw data. For example, smaller
13 identifiers can be more easily sorted and searched using standard methods. Thus,
14 hashing generally yields greater benefits when smaller hash values are used.

15 Unfortunately, there is a point at which hash values become too small and
16 begin to lose the desirable quality of uniquely representing a large mass of data
17 items. That is, as the size of hash values decreases, it is increasingly likely that
18 more than one distinct raw data can be mapped into the same hash value, an
19 occurrence referred to as “collision”. Mathematically, for A alphabets of each
20 hash digit and a hash value length l , an upper bound of all possible hash values is
21 A^l . If the number of distinct raw data are larger than this upper bound, collision
22 will occur.

23 Accordingly, another property of a good hash function is to minimize the
24 probability of collision. However, if considerable gain in the length of the hash
25 values can be achieved, it is sometimes justified to tolerate collision. The length

1 of the hash value is thus a trade off with probability of collision. A good hash
2 function should minimize both the probability of collision and the length of the
3 hash values. This is a concern for design of both hash functions in compilers and
4 message authentication codes (MACs) in cryptographic applications.

5 Good hash functions have long existed for many kinds of digital data.
6 These functions have good characteristics and are well understood. The idea of a
7 hash function for image database management is very useful and potentially can
8 be used in identifying images for data retrieval and copyrights protection.
9 Unfortunately, while there are many good existing functions, digital images
10 present a unique set of challenges not experienced in other digital data, primarily
11 due to the unique fact that images are subject to evaluation by human observers. A
12 slight cropping or shifting of an image does not make much difference to the
13 human eye, but such changes appear very differently in the digital domain. Thus,
14 when using conventional hashing functions, a shifted version of an image
15 generates a very different hash value as compared to that of the original image,
16 even though the images are essentially identical in appearance. Another example
17 is the deletion of one line from an image. Most people will not recognize this
18 deletion in the image itself, yet the digital data is altered significantly if viewed in
19 the data domain.

20 Human eyes are rather tolerant of certain changes in images. For instance,
21 human eyes are much less sensitive to high frequency components of an image
22 than low frequency components. In addition, the average (i.e., DC component) is
23 interpreted by our eyes as brightness of an image and it can be changed within a
24 range and cause only minimal visible difference to the observer. Our eyes would
25 also be unable to catch small geometric deformation in most images.

1 Many of these characteristics of the human visual system can be used
2 advantageously in the delivery and presentation of digital images. For instance,
3 such characteristics enable compression schemes, like JPEG, to compress images
4 with good results, even though some of the image data may be lost or go unused.
5 There are many image restoration/enhancement algorithms available today that are
6 specially tuned to the human visual system. Commercial photo editing systems
7 often include such algorithms.

8 At the same time, these characteristics of the human visual system can be
9 exploited for illegal or unscrupulous purposes. For example, a pirate may use
10 advanced image processing techniques to remove copyright notices or embedded
11 watermarks from an image without visually altering the image. Such malicious
12 changes to the image are referred to as “attacks”, and result in changes at the data
13 domain. Unfortunately, the user is unable to perceive these changes, allowing the
14 pirate to successfully distribute unauthorized copies in an unlawful manner.
15 Traditional hash functions are of little help because the original image and pirated
16 copy hash to very different hash values, even though the images appear the same.

17 Accordingly, there is a need for a hash function for digital images that
18 allows slight changes to the image which are tolerable or undetectable to the
19 human eye, yet do not result in a different hash value. For an image hash function
20 to be useful, it should accommodate the characteristics of the human visual system
21 and withstand various image manipulation processes common to today’s digital
22 image processing. A good image hash function should generate the same unique
23 identifier even though some forms of attacks have been done to the original image,
24 given that the altered image is reasonably similar to a human observer when
25 comparing with the original image. However, if the modified image is visually

1 different or the attacks cause irritation to the observers, the hash function should
2 recognize such degree of changes and produce a different hash value from the
3 original image.

4 5 **SUMMARY**

6 This invention concerns a system and method for hashing digital images in
7 a way that allows modest changes to an image, which may or may not be
8 detectable to the human eye, yet does not result in different hash values for the
9 original and modified images.

10 According to one implementation, a system stores original images in a
11 database. An image hashing unit hashes individual images to produce hash values
12 that uniquely represent the images. The image hashing unit implements a hashing
13 function H , which takes an image I and an optional secret random string as input,
14 and outputs a hash value X according to the following properties:

- 15
16 1. For any image I_i , the hash of the image, $H(I_i)$, is approximately random
17 among binary strings of equal length.
- 18 2. For two distinct images, I_1 and I_2 , the hash value of the first image,
19 $H(I_1)$, is approximately independent of the hash value of the second
20 image, $H(I_2)$, in that given $H(I_1)$, one cannot predict $H(I_2)$ without
21 knowing a secret key used to produce $H(I_1)$.
- 22 3. If two images I_1 and I_2 are visually the same or similar, the hash value
23 of the first image, $H(I_1)$, should equal the hash value of the second
24 image, $H(I_2)$.

1 The hash value is stored in an image hash table and is associated via the
2 table with the original image I from which the hash is computed. This image hash
3 table can be used to index the image storage.

4 The processing system also has a watermark encoder to watermark
5 individual images. The watermark encoder computes a watermark based on the
6 hash value X and a secret W . Using both values effectively produces unique
7 secrets for each individual image. Thus, even if the global watermark secret is
8 discovered, the attacker still needs the hash value of each image to successfully
9 attack the image. As a result, the system is resistant to BORE (Break Once, Run
10 Everywhere) attacks, thereby providing additional safeguards to the images.

11 The watermark encoder encodes the watermark into the original image I to
12 produce a watermarked image I' . The system may store and/or distribute the
13 watermarked image.

14 According to an aspect of this invention, the system can be configured to
15 search over the Internet to detect pirated copies. The system randomly collects
16 images from remote Web sites and hashes the images using the same hashing
17 function H . The system then compares the image hashes to hashes of the original
18 images. If the hashes match, the collected image is suspected as being a copy of
19 the original.

20 21 **BRIEF DESCRIPTION OF THE DRAWINGS**

22 The same numbers are used throughout the drawings to reference like
23 elements and features.

1 Fig. 1 is a block diagram of an image distribution system in which a content
2 producer/provider hashes and watermarks digital images and subsequently
3 distributes them to a client over a network.

4 Fig. 2 is a functional block diagram of an image hash unit implemented at
5 the content producer/provider of Fig. 1 to hash the digital images.

6 Fig. 3 is a diagrammatic illustration of a process of dividing an image
7 transform into multiple non-overlapping tiles.

8 Fig. 4 is a diagrammatic illustration of a process of dividing an image
9 transform into multiple overlapping tiles.

10 Fig. 5 is a diagrammatic illustration of quantization points to demonstrate a
11 process of rounding tile averages to one of the points.

12 Fig. 6 is a flow diagram showing a method for distributing watermarked
13 digital images over a network and through surveillance, detecting pirated versions
14 of the digital images using a hash compare operation.

15 16 **DETAILED DESCRIPTION**

17 This invention is described below as a technique for hashing digital images.
18 Thus, the described hashing techniques are particularly tailored to accommodate
19 characteristics of the human visual system and withstand various image
20 manipulation processes common to today's digital image processing. However,
21 the invention is not limited in its application to digital images. Rather, the
22 described techniques can also be applied to other sampled or digitized media
23 streams such as digitized audio streams.

24 The described hashing techniques generate the same unique identifier even
25 though some forms of attacks have been done to the original image, given that the

1 altered image is reasonably similar to a human observer when comparing the
2 altered image with the original image. However, if the altered image is visually
3 different or the attacks cause irritation to the observers, the hash function
4 recognizes such degree of changes and produces a different hash value from the
5 original image.

6 The hash function implemented by various systems and methods described
7 herein is denoted as H . Given an input image I , the hash function H produces a
8 short binary string X , as follows:

$$H(I) = X$$

11
12 The hash function H has the following properties:

- 13
- 14 4. For any image I_i , the hash of the image, $H(I_i)$, is approximately random
15 among binary strings of equal length.
 - 16 5. For two distinct images, I_1 and I_2 , the hash value of the first image,
17 $H(I_1)$, is approximately independent of the hash value of the second
18 image, $H(I_2)$, in that given $H(I_1)$, one cannot predict $H(I_2)$ without
19 knowing a secret key used to produce $H(I_1)$.
 - 20 6. If two images I_1 and I_2 are visually the same or similar, the hash value
21 of the first image, $H(I_1)$, should equal the hash value of the second
22 image, $H(I_2)$.

23
24 A special case of the third property is where an original image is attacked to
25 remove the watermark or copyright notice. In this case, suppose the original

1 image I_O is modified to include a watermark, thus producing a watermarked image
2 I_{WM} . Using property three, the images are visually identical and hence, $H(I_O) =$
3 $H(I_{WM})$. Now, suppose that the watermarked image is attacked using digital image
4 processing techniques to remove the watermark and produce a pirate image I_P ,
5 which is visually identical to the original image I_O and the watermarked image
6 I_{WM} . In this case, the hash values are also the same, i.e., $H(I_{WM}) = H(I_P)$.

7 One exemplary implementation of the hashing function H is described
8 below in more detail. In addition, exemplary implementations of the hashing
9 technique in various systems and methods are described below, beginning with an
10 architecture for electronic distribution of digital images over a network, such as
11 the Internet.

12 13 System Architecture

14 Fig. 1 shows an image distribution system 20 having a content
15 producer/provider 22 that produces digital images and/or distributes the digital
16 images over a network 24 to a client 26. The content producer/provider 22 has an
17 image storage 30 to store digital images, a processing system 32 to process the
18 images prior to distribution, and a distribution server 34 to distribute the images
19 over the network 24 (e.g., Internet, LAN, WAN, etc.). The server 34 may be
20 further configured to compress and/or encrypt the images using conventional
21 compression and encryption techniques prior to distributing the content over the
22 network 24.

23 The processing system 32 has an image hashing unit 40 that hashes
24 individual images to produce hash values that uniquely represent the images. The
25 image hashing unit 40 implements the hashing function H , which takes an image I

1 as input, and outputs a hash value X according to the properties described above.
2 The hash value is stored in an image hash table 44 in storage 30 and is associated
3 via the table 44 with the original image I from which the hash is computed. This
4 image hash table 44 can be used to index the image storage 30.

5 The processing system 32 also has a watermark encoder 42 to watermark
6 individual images. A watermark is an array of bits generated using known
7 cryptographic techniques and embedded into a digital image, without affecting the
8 appearance of the image. The watermark encoder 42 receives the hash value X ,
9 and computes a watermark based, in part, on the hash value X and a secret W . The
10 watermark encoder 42 encodes the watermark into the original image I to produce
11 a watermarked image I' . The system 32 may store the watermarked image I' in
12 the image storage 30 and/or passed to the distribution server 34 for distribution
13 over the network 24 to the client 26.

14 An advantage of computing the watermark based on the hash value X is that
15 it adds security on a per image basis. Normally, a single watermark based on the
16 watermark secret W is globally applied to all images in the storage 30. In contrast,
17 image hash unit creates separate and distinct hash values X for each of the images.
18 The watermark encoder 42 then uses these values in conjunction with the
19 watermark secret W to effectively produce unique secrets for each individual
20 image. Thus, even if the watermark secret is discovered, the attacker still needs
21 the hash value of each image to successfully attack the image. As a result, the
22 system is resistant to BORE (Break Once, Run Everywhere) attacks, thereby
23 providing additional safeguards to the images.

24 It is noted that the image hashing unit 40 and watermark encoder 42 may be
25 implemented in software or firmware. These components may be constructed as

1 part of a processing system, incorporated into other applications or an operating
2 system, or formed as separate standalone modules. The content producer/provider
3 22 may be implemented in many ways, including as one or more server computers
4 configured to store, process, and distribute digital images.

5 The client 26 is equipped with a processor 50, a memory 52, and one or
6 more media output devices 54. The processor 50 runs various tools to process the
7 digital images, such as tools to decompress the images, decrypt the data, and/or
8 apply controls (size, rotation, etc.). The memory 52 stores an operating system 56,
9 such as a Windows brand operating system from Microsoft Corporation, which
10 executes on the processor. The client 26 may be embodied in a many different
11 ways, including a computer, a handheld entertainment device, a set-top box, a
12 television, and so forth.

13 The operating system 56, or any trusted software or hardware on the client
14 machine, may implement a client-side watermark detector 58 to detect the
15 watermark in the digital images. If the watermarks are present, the client is
16 assured that the content is original and can be played. Absence of the watermark
17 indicates that the image is a pirated copy of the original. The operating system 56
18 and/or processor 50 may be configured to enforce certain rules imposed by the
19 content producer/provider (or copyright owner). For instance, the operating
20 system and/or processor may be configured to reject fake or copied images that do
21 not possess a watermark.

22 23 **Image Hash Unit**

24 Fig. 2 shows the image hash unit 40 in more detail. The image hash unit 40
25 includes a random linear coder 70, an image transformer 72, a tile creation and

1 averaging module 74, a randomized rounding module 76, an intermediate hashing
2 module 78, and an error correction module 80. These components are preferably
3 implemented in software, although some or all of the components may be
4 implemented in firmware or hardware.

5 The random linear coder 70 selects a random linear code C that is used later
6 in the processing by the error correction module. The selection is performed once
7 during initialization and the linear code C is used for all images. The linear code
8 C has three selectable parameters n' , k' , and d' , where n' represents a length of a
9 random string, k' represents a length of the original message, and k' represents a
10 Hamming distance. The linear code C has the following properties:

$$C = \{B : AB = 0 \text{ mod } 2\}$$

11
12
13
14 where A is an $m' \times n'$ matrix (m' is computed from n' , k' , and d') in which each
15 entry is chosen randomly from a set of values $\{0, 1\}$. B is an $n' \times 1$ matrix
16 containing an n' -bit array generated by the pseudo-random number generator. As
17 an example, the set of parameters (n', k', d') equal (32, 15, 5).

18 The image transformer 72 receives an original image I and computes a
19 transformation of the image using a transform function f , where $f = (f_1, f_2, \dots, f_n)$.
20 The transformer 72 may use one of many conventional transforms, such as a
21 Fourier transform, a wavelet transform, and a DCT (Discrete Cosine transform)
22 transform. As one exemplary implementation, the image transformer 72 uses a
23 wavelet transform to decompose the image into three levels or subbands: a coarse
24 subband, a high-low subband, and a low -high subband.
25

1 The tile creation and averaging module 74 randomly divides the image
2 transform image into multiple tiles t , where each tile contains data for multiple
3 pixels in the image. Two possible techniques for constructing the tiles are (1)
4 forming non-overlapping rectangular tiles and (2) creating overlapping rectangular
5 tiles.

6 Fig. 3 illustrates the process of forming non-overlapping tiles. Given image
7 transform 90, the tile creation module 74 randomly picks a point P_1 somewhere
8 between one-third and two-thirds of the base and divides the image transform 90
9 into side-by-side rectangles. The module 74 then randomly selects a point Q_1
10 somewhere between one-third to two-thirds of the height and divides the left-side
11 rectangle. One can use any suitable distribution that splits the image into
12 approximately equal portions here. Similarly, the module 74 randomly selects a
13 point Q_2 somewhere between one-third to two-thirds of the height and divides the
14 right-side rectangle. This process is repeated for each of the subrectangles until a
15 predetermined number of tiles is created.

16 Fig. 4 illustrates the process of forming overlapping tiles. Given the image
17 transform 90, the tile creation module 74 randomly draws a series of rectangles
18 over the space, until the predetermined number of tiles is created.

19 In the continuing exemplary implementation in which a wavelet transform
20 is used, each subband—coarse, low-high, and high-low—is divided into 64 small
21 rectangles with random sizes. The coarse subband is divided using the non-
22 overlapping process of Fig. 3, whereas the high and high-high subbands are
23 divided using the overlapping process of Fig. 4. This produces 192 tiles for an
24 image.
25

1 After tile creation, the tile creation and averaging module 74 computes an
2 average of each tile t . In the continuing example involving a wavelet transform,
3 suppose that a tile of the transformed image has data for pixels g_1, g_2, \dots, g_N . The
4 tile creation and averaging module 74 produces an average μ for each tile t , as
5 follows:

$$u_i = \frac{\sum_{j=1}^N g_j}{N}$$

8
9 In high and high-high subbands, the average may be zero. In those cases,
10 the variance is computed instead of the average, as follows:

$$v^2 = \frac{\sum_{j=1}^N g_j^2}{N}$$

14 For the coarse subband, the module 74 quantizes the averages of the tiles
15 into eight (an example value) distinct levels based on an absolute value of the tile
16 average. The maximum value for the coarse subband is 2040 (i.e., 255×8) and this
17 value is divided by eight to produce eight different quantization levels. A total of
18 126 values are computed for the 64 random tiles and their combinations.

19 As for the high-low and low -high subbands, where the average values are
20 guaranteed to be zero, the variances are computed and fitted by an exponential
21 distribution so that the values fall into approximately eight different levels. To
22 reduce the effect from exclusion/inclusion of an edge in the tile due to shifting of
23 an image, a window function is used around the tile to reduce the effect of edges.
24
25

1 The tile creation and averaging module 74 outputs an average vector $\mu =$
2 $(\mu_1, \mu_2, \dots, \mu_t)$ having averages for the t tiles in the image.

3 With continuing reference to Fig. 2, the randomized rounding module 76
4 receives the average vector and computes, for each subband, the average of all the
5 tiles' variances. Using this average, the rounding module 76 creates an
6 exponential distribution and generates eight distinct quantization levels based on
7 this distribution. Each quantization level has a probability mass of one-eighth,
8 meaning that for a random tile the results from rounding will be uniformly
9 distributed across the quantization levels. The quantization levels are represented
10 as A_0 to A_7 . The rounding module 76 rounds each of the averages μ_i for each tile t
11 to one of the eight quantization levels.

12 Fig. 5 illustrates the rounding process. Suppose that an average μ_i falls
13 between quantization levels A_2 and A_3 . The rounding module 76 tends to favor
14 rounding the average μ_i toward the closer of the two quantization levels, which in
15 this case is level A_2 . But, the rounding module 76 also introduces some
16 randomness to make it more difficult for an attacker to predict the outcome of the
17 rounding. The randomness essentially imposes a coin flip strategy in which the
18 mathematical expectation of the outcome after the rounding is equal to the original
19 value of the quantity being rounded. Stated alternatively, the mathematical
20 expectation of the outcome should be equal to a continuous function of the value
21 being rounded. Also one may use a buffered rounding strategy where the given
22 quantity (μ_i) is rounded to the nearest number A_2 or A_3 , if the distance to the
23 nearest number is smaller than some pre-determined bound. As a result, the
24 rounding module 76 effectively rounds the average μ_i toward one of the two
25 quantization levels according to a coin flip that is biased slightly toward rounding

1 to the nearest of the two quantization levels. It is further noted that the
2 quantization levels are generated in the first place with a goal of making the
3 expected value of the rounding to be equal to the original value, and any small
4 changes in the given quantity will result in small changes in the expected value of
5 the final rounded output.

6 Mathematically, let ρ represent a distance parameter involving μ_i and the
7 two nearest quantization levels A_2 and A_3 . By this we mean

$$\rho = (\mu_i - A_2) / (A_3 - A_2); \text{ or}$$

9
10 Now we flip a coin which has bias ρ of getting heads and $1-\rho$ for tails.
11 Then if we get a head we round μ_i to A_3 . Otherwise, we round it to A_2 . The
12 distance parameter ρ generates a bias towards the closer value while the random
13 number r provides some randomness to make it more difficult for the attacker to
14 predict the outcome.

15 The rounded values are mapped into k -bit binary strings q_1, q_2, \dots, q_t , one
16 string for each tile t . In our continuing example, the rounded values are mapped
17 into 3-bit binary strings representative of the quantization points $A_0 \dots A_7$. A
18 value μ_i that rounds to A_0 is mapped to binary string "000", a value μ_i that rounds
19 to A_1 is mapped to "001", a value μ_i that rounds to A_2 is mapped to "010", and so
20 on. As noted above, the points $A_0 \dots A_7$ are chosen so that for a random tile, the
21 rounding steps yields uniformly distributed 3-bit binary strings. The k -bit binary
22 strings are appended together to form a composite value Q , which is output by the
23 randomized rounding module 76.

24 The rounding sub-process provides particular advantages for the image
25 hashing process. The rounded values are used instead of the precise averages in

1 later computation of the hash values. In this manner, slight modifications can be
2 made to an image without changing the hash value for the image. That is, an
3 attacker can make minor changes, such as removing a watermark, that modify the
4 averages μ_i for some or all of the tiles. Without rounding and the subsequent
5 error-correction sections, these changes would result in a different hash value.
6 However, with the rounding process and the error correction, these changes do not
7 carry through to the resulting rounded values and hence, the hash value for the two
8 images remains the same.

9 The intermediate hash module 78 receives the composite value Q and
10 produces an intermediate hash IH with the following properties:

- 11
12 1. For two visually distinct images I_1 and I_2 , the intermediate hash values
13 differ approximately 60% of the time.
- 14 2 For two visually similar images I_1 and I_2 , the intermediate hash values
15 agree in all but approximately 20% of the time.

16
17 The above numbers (60% 20%) are indicative of the exemplary
18 implementation and can vary depending on the characteristics of the digitized
19 stream.

20 In the continuing example, the intermediate hashing module 78 implements
21 a first order Reed-Muller error correction code decoder. Such decoders are well
22 known and other error correcting code decoders may be used (See, e.g., NJA
23 Sloane and Mc Williams, "Theory of Error Correcting Codes", North Holland).
24 The Reed-Muller decoder (or other suitable decoder) is modified, however, to
25

1 work with a distance function we call an exponential pseudo-norm. Given a
2 vector $v = (v_1, v_2, \dots, v_l)$, the pseudo-random norm is:

$$3 \quad 4 \quad 5 \quad 6 \quad norm(v) = \sum_{i=1}^l C^{|v_i|}$$

7 It is noted that the image hash unit 40 does not employ a complementary
8 encoder, but only the decoder.

9 The error correction module 80 receives the intermediate hash IH and
10 reduces the hash size and number of error occurrences. In our continuing
11 example, the error correction module 80 extracts a subset of bits from the
12 intermediate hash IH . The subset of bits is chosen so that approximately one-half
13 of the bits are extracted from the coarse subbands and one-fourth of the bits are
14 extracted from each of the two high frequency subbands. There are hundreds of
15 bits in the intermediate hash IH and the extracted subset of bits typically numbers
16 less than one hundred.

17 From the subset, the error correction module 80 further extracts a reduced
18 set of bits, such as 32 bits. This reduced set of bits is then processed using a list-
19 decoding process into a small list $\{X_1, X_2, \dots, X_r\}$, where r is small. List-decoding
20 is well known. For a very brief discussion on list-decoding, the reader is directed
21 to the last section of L.A. Levine, "One-Way Functions and Pseudo-Random
22 Generators", *Combinatorica* 7, 1987, pgs. 357-363, and to P. Elias, "Personal
23 Communication to L.A. Levin", 1988. Also see the following references:

- 24 • Sudan, Madhu; *Proceedings of the 37th Annual IEEE Symposium on*
25 *Foundations of Computer Science*; "Maximum Likelihood Decoding

of Reed Solomon Codes”; 1996 (a more recent version of this paper, entitled “Decoding Reed Solomon Codes Beyond the Error-Correction Bound”, is available by request from MIT Laboratory of Computer Sciences, and is available on the Internet at the time of this writing at [“http://theory.lcs.mit.edu/~madhu/papers/reedsolomon.ps”](http://theory.lcs.mit.edu/~madhu/papers/reedsolomon.ps).

- *Journal of Complexity* , Special issue dedicated to Shmuel Winograd, 13(1): 180-193, March 1997

One word is selected from this list using image parameters/semantics to produce a final hash value X . Specifically, a word is selected using a maximum likelihood method. In the example, the resultant hash value is 32 bits. However, if this value results in a higher probability of collision, a longer hash value may be obtained by running the process twice to produce two 32-bit values or by increasing the parameters of the linear coding unit 70 to produce larger encoded messages.

Exemplary Contexts

The image hashing process described above, and implemented in the image hashing unit 40, can be used in many ways and in a number of contexts. For instance, the image hashing process can be used as an indexing system for a large database of images. In this context, the image hashes X are stored in an indexing table 44 (Fig. 1) and used to rapidly index the associated images in the image storage 30.

1 Another exemplary context is to use the image hashing process as a way to
2 police search over the Internet to detect pirated copies. Generally, this is done by
3 randomly collecting images, hashing them, and comparing the image hashes to
4 hashes of the original images. If the hashes match, the collected image is
5 suspected as being a copy of the original.

6 Fig. 6 illustrates a detailed process of distributing watermarked digital
7 images and through surveillance, detecting pirated versions of the digital images
8 using the image hash process. The process is implemented primarily in software,
9 although aspects may be implemented using hardware and firmware. The process
10 is further described with reference to Fig. 1.

11 At step 100, the processing system 32 of the content producer/provider 22
12 retrieves an image from the image store 30 and computes an image hash X
13 associated with the image. The image hash X is stored in the image hash table 44
14 and associated with the original image. The processing system 32 then
15 watermarks the image using the image hash X and a secret key W to produce the
16 watermark (step 102). This combination of secrets makes the watermark unique to
17 each image, rather than global to all images. The watermarked images may
18 optionally be stored in the image storage 30.

19 At step 104, the distribution server 34 distributes the watermarked image I'
20 over the network 24 to a client 26. In this case, suppose the client is a pirate who
21 intends to attack the image and remove the watermark (step 106). Through the
22 attacks, the pirate is able to produce a pirated version of the image that is visually
23 identical or very similar, but without the watermark (step 108). The pirate then
24 redistributes the pirated version for illicit gain (step 110).
25

1 Through standard surveillance practices, the original content
2 producer/provider 22 routinely and randomly gathers images from various Web
3 sites. In a routine sweep, the content producer/provider 22 collects the pirated
4 version along with other images (step 112). The content producer/provider 22
5 uses the image hash unit 40 to compute image hashes of each collected image
6 (step 114). The content producer/provider 22 then compares each image hash of
7 the collected images with image hashes stored in the image hash table 44 to
8 evaluate whether any match occurs (step 116). If the image hash of a collected
9 image matches a stored image hash (i.e., the “yes” branch from step 118), the
10 image is detected as potentially being a pirated version (step 120). Conversely, if
11 no match occurs, the collected versions are not considered duplicate or altered
12 versions of the original images (step 122).

13 14 **Conclusion**

15 Although the invention has been described in language specific to structural
16 features and/or methodological steps, it is to be understood that the invention
17 defined in the appended claims is not necessarily limited to the specific features or
18 steps described. Rather, the specific features and steps are disclosed as preferred
19 forms of implementing the claimed invention.
20
21
22
23
24
25